Chapter 2
Network Architecture and Design
NET3011 – 16W

Campus Network

A campus network describes the portion of an enterprise infra-structure that interconnects end devices such as computers, laptops, and wireless access points to services such as intranet resources or the Internet." p. 9, previous version of FLG

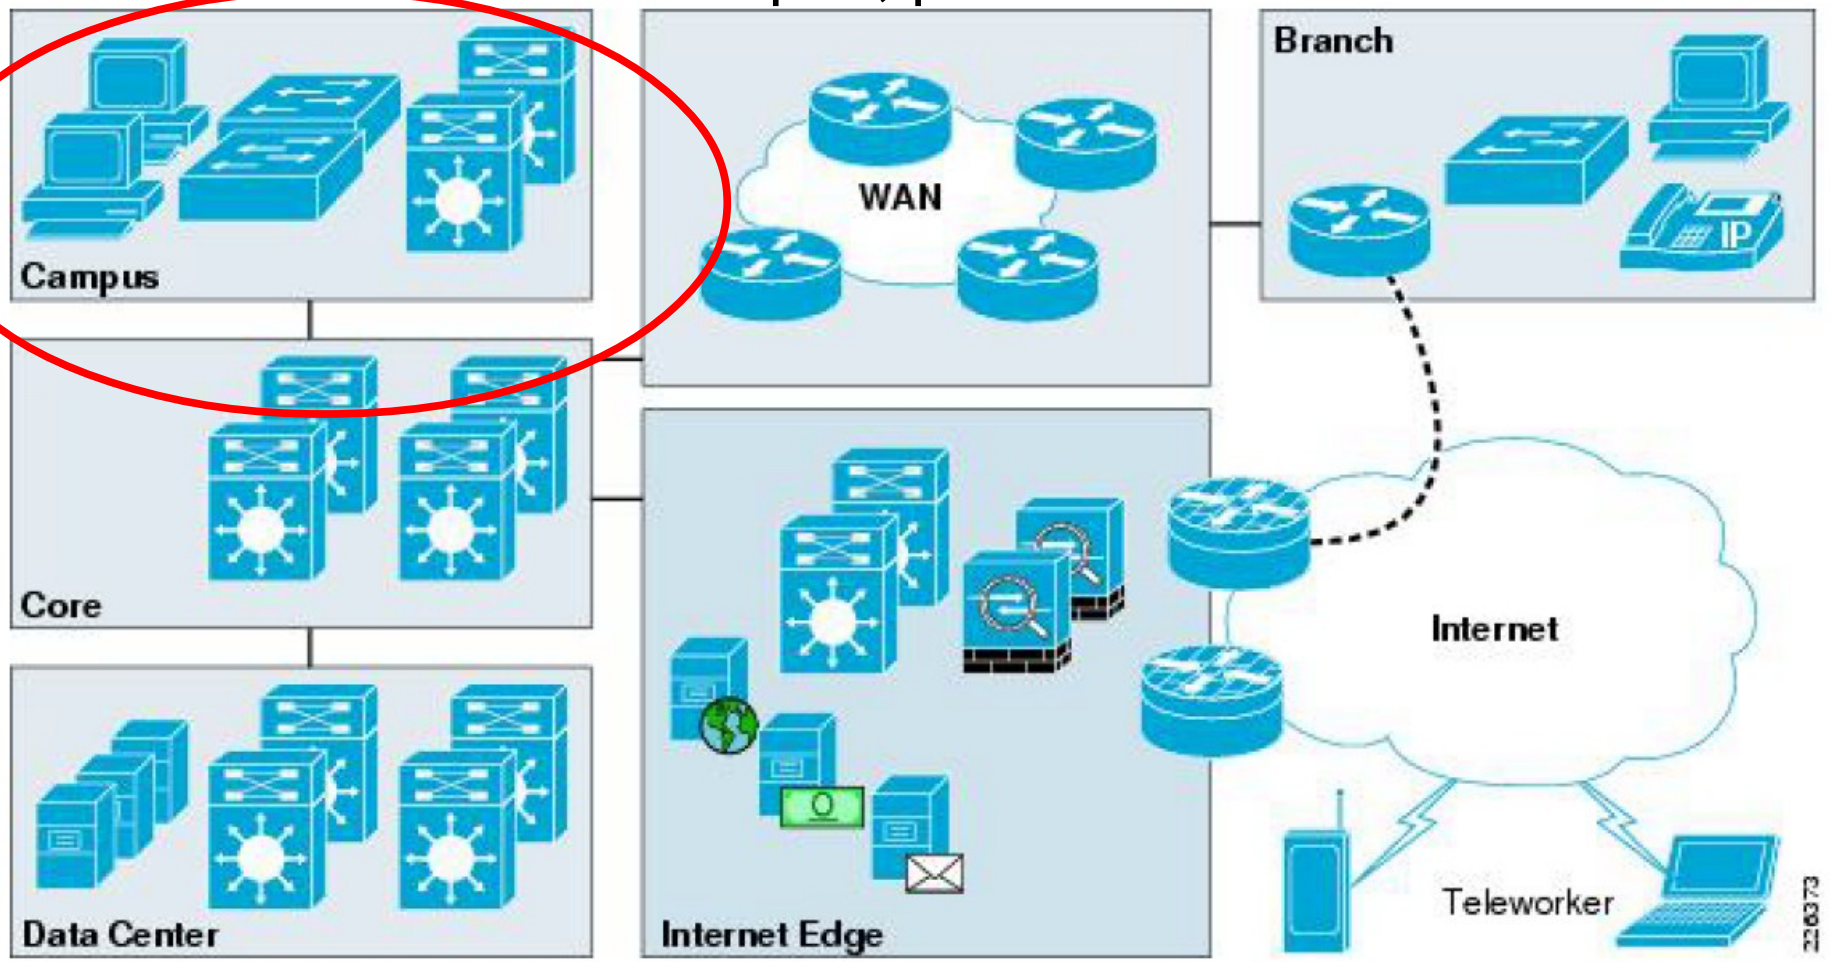

Flat vs Hierarchical Networks

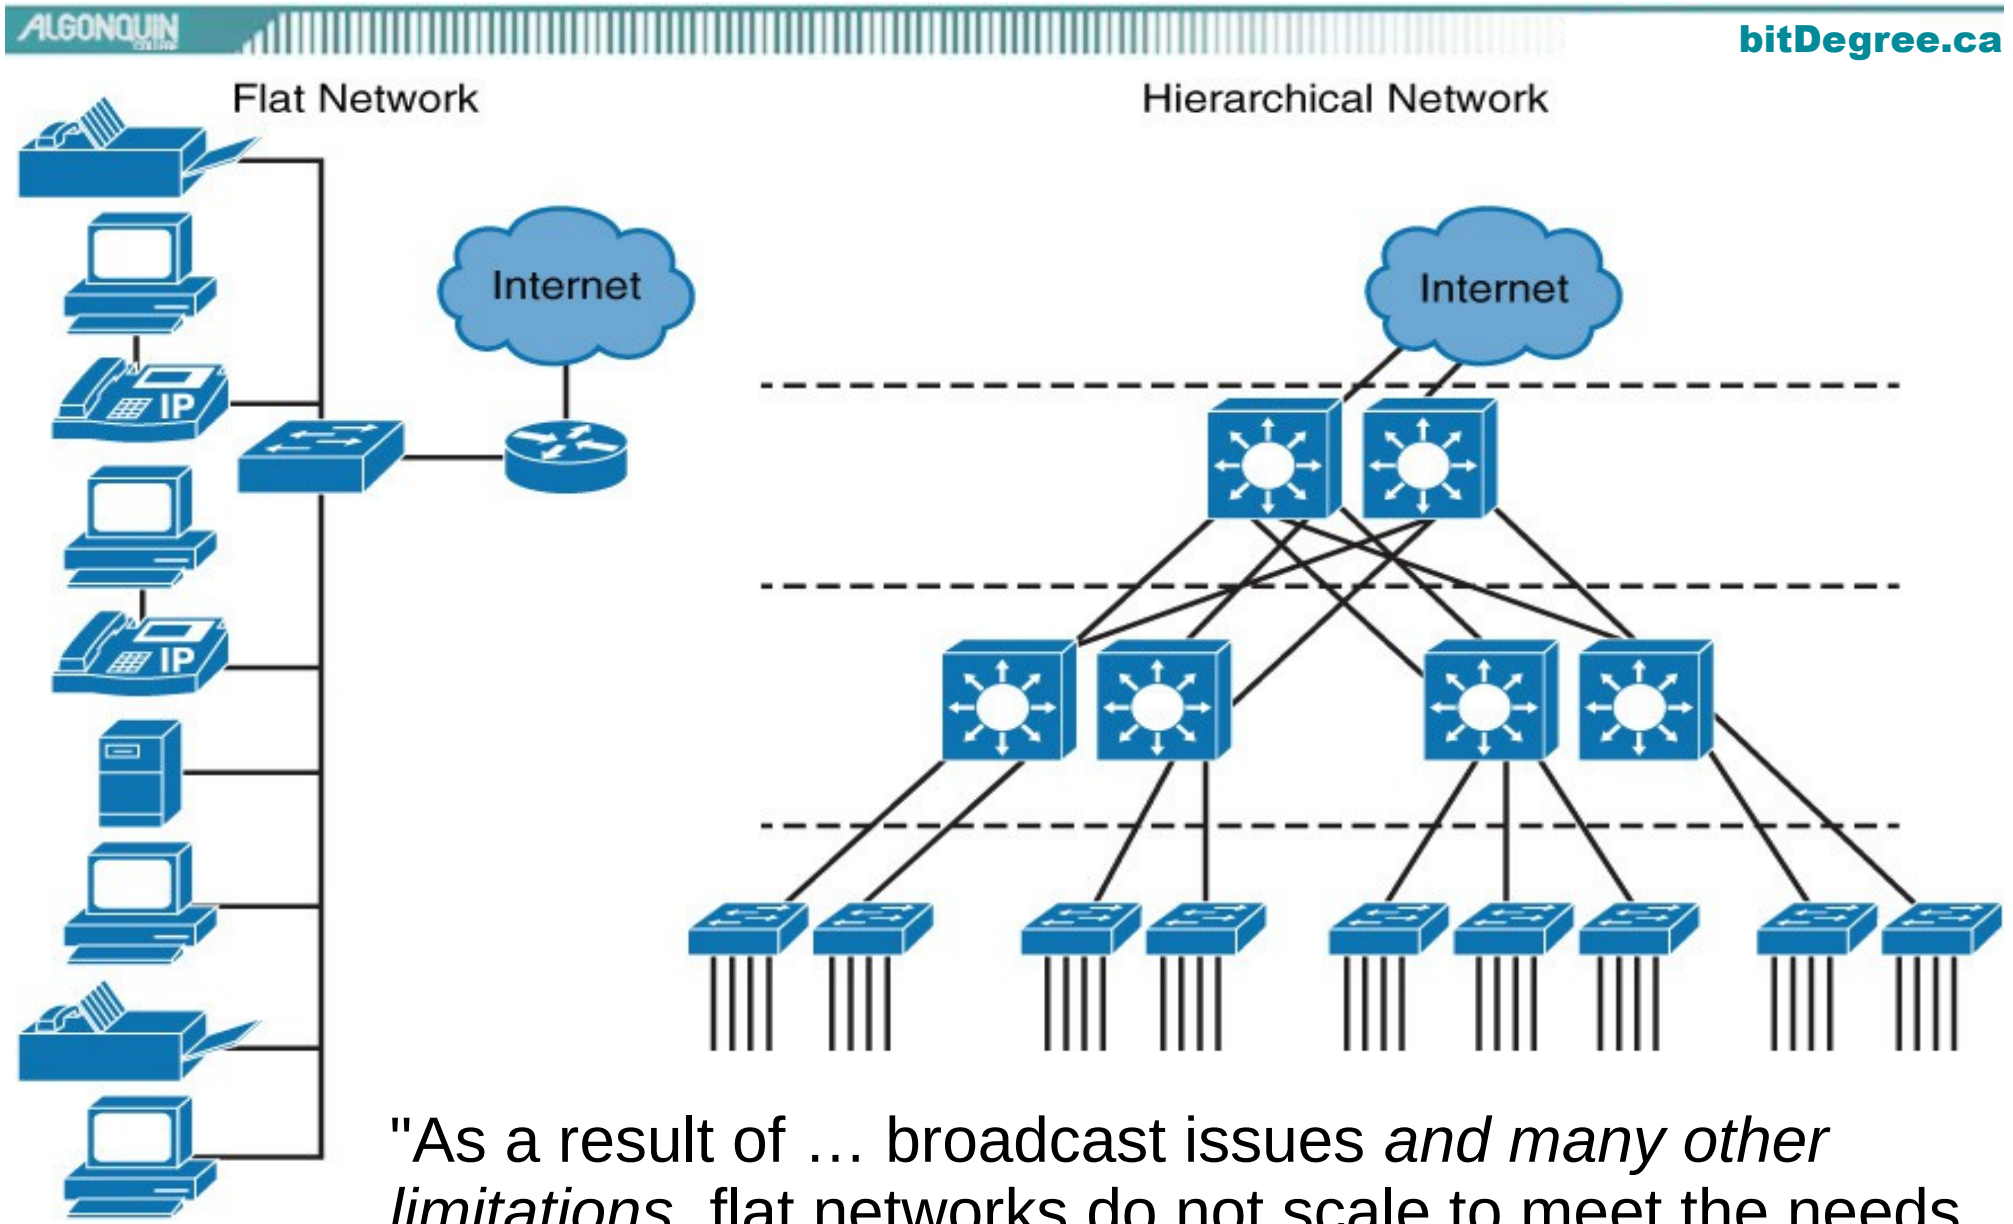

"As a result of … broadcast issues *and many other limitations*, flat networks do not scale to meet the needs of enterprise networks." p. 11 FLG

So what's wrong with a flat network?

Plenty!

- Single point of failure for Layer 2 and Layer 3 devices
- Multilayer switching functionality may be under-utilized if a multilayer switch is deployed as a simple physical replacement for a traditional router or switch without any network re-design
- Increased Spanning Tree complexity
- Servers not centrally located
- Large broadcast domains
- May have high latency
- Difficult to troubleshoot

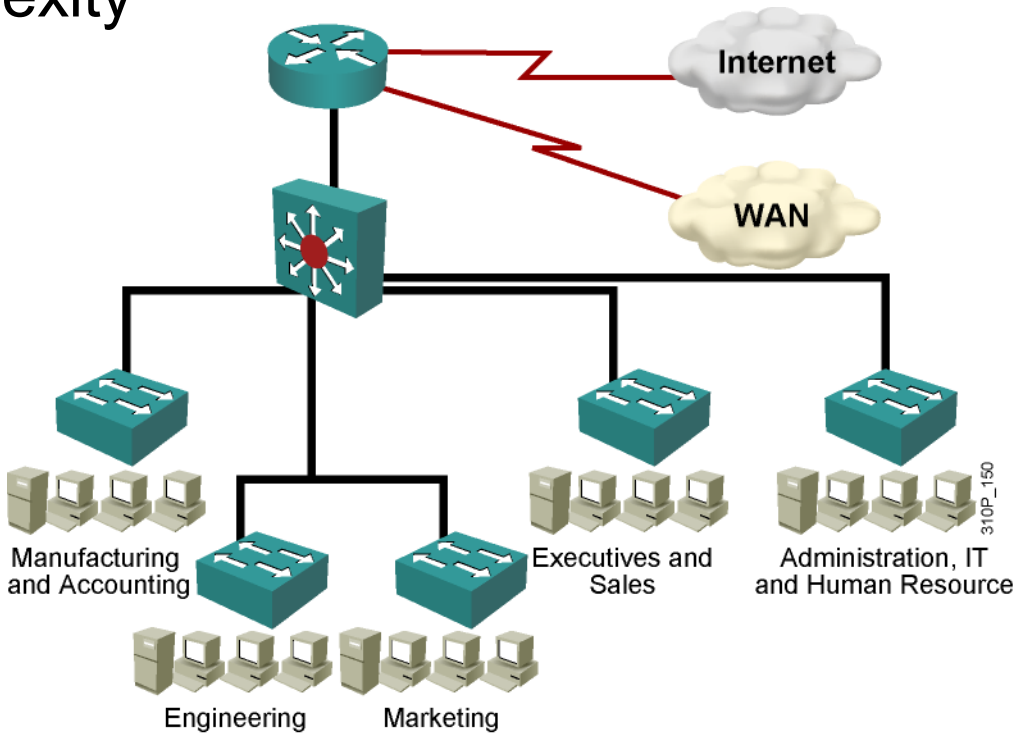

Hierarchical Network Model

Allows flexibility in network design and facilitates ease of implementation and troubleshooting.

Modular!!!

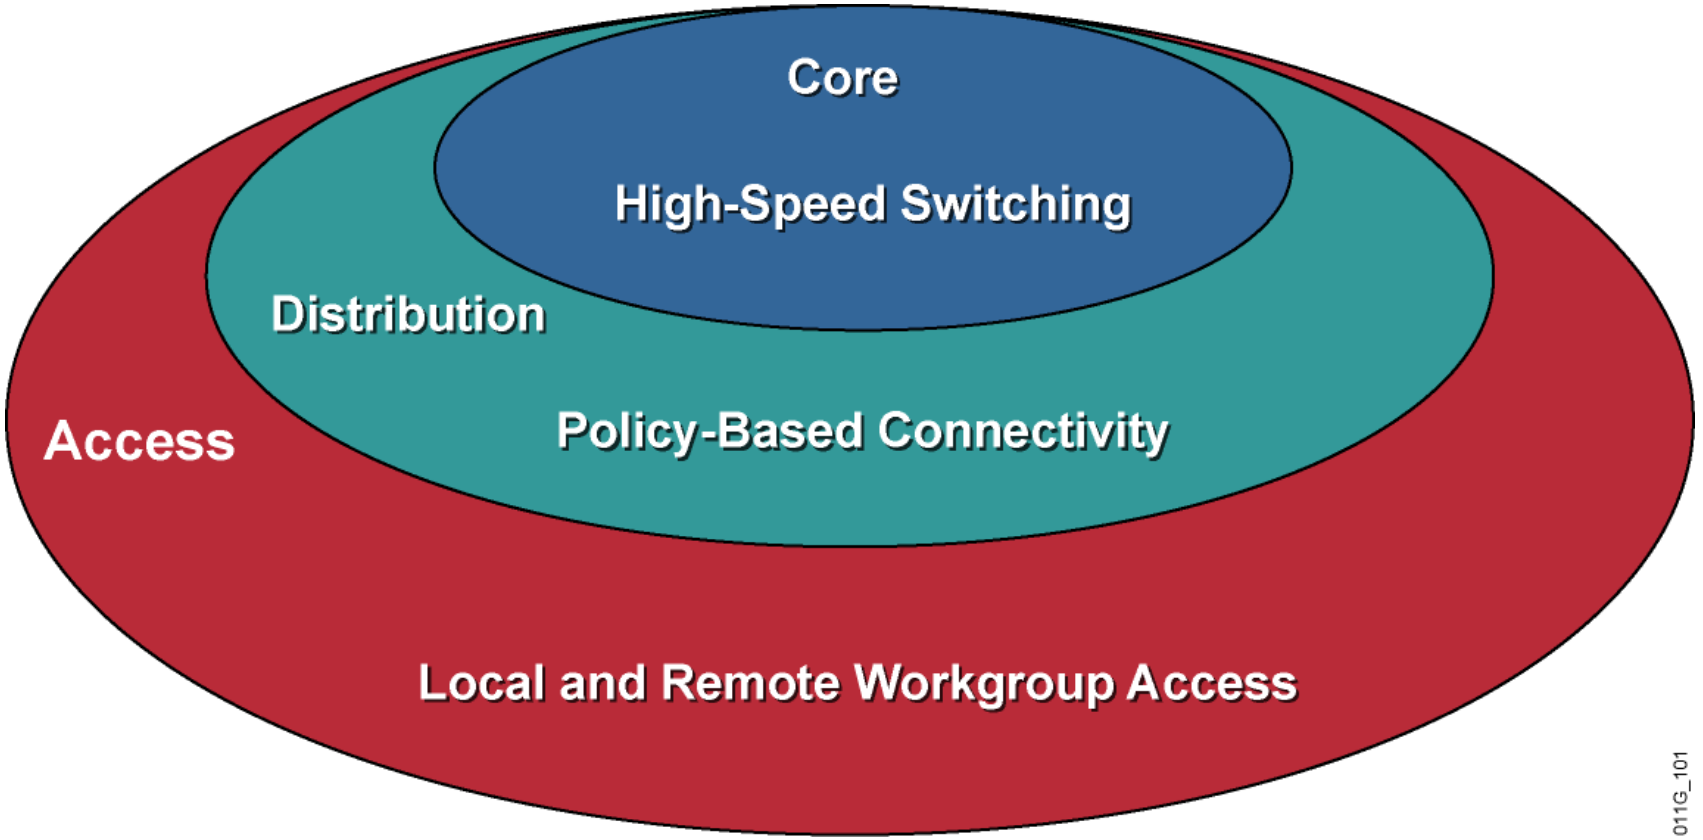

Access Layer – in context

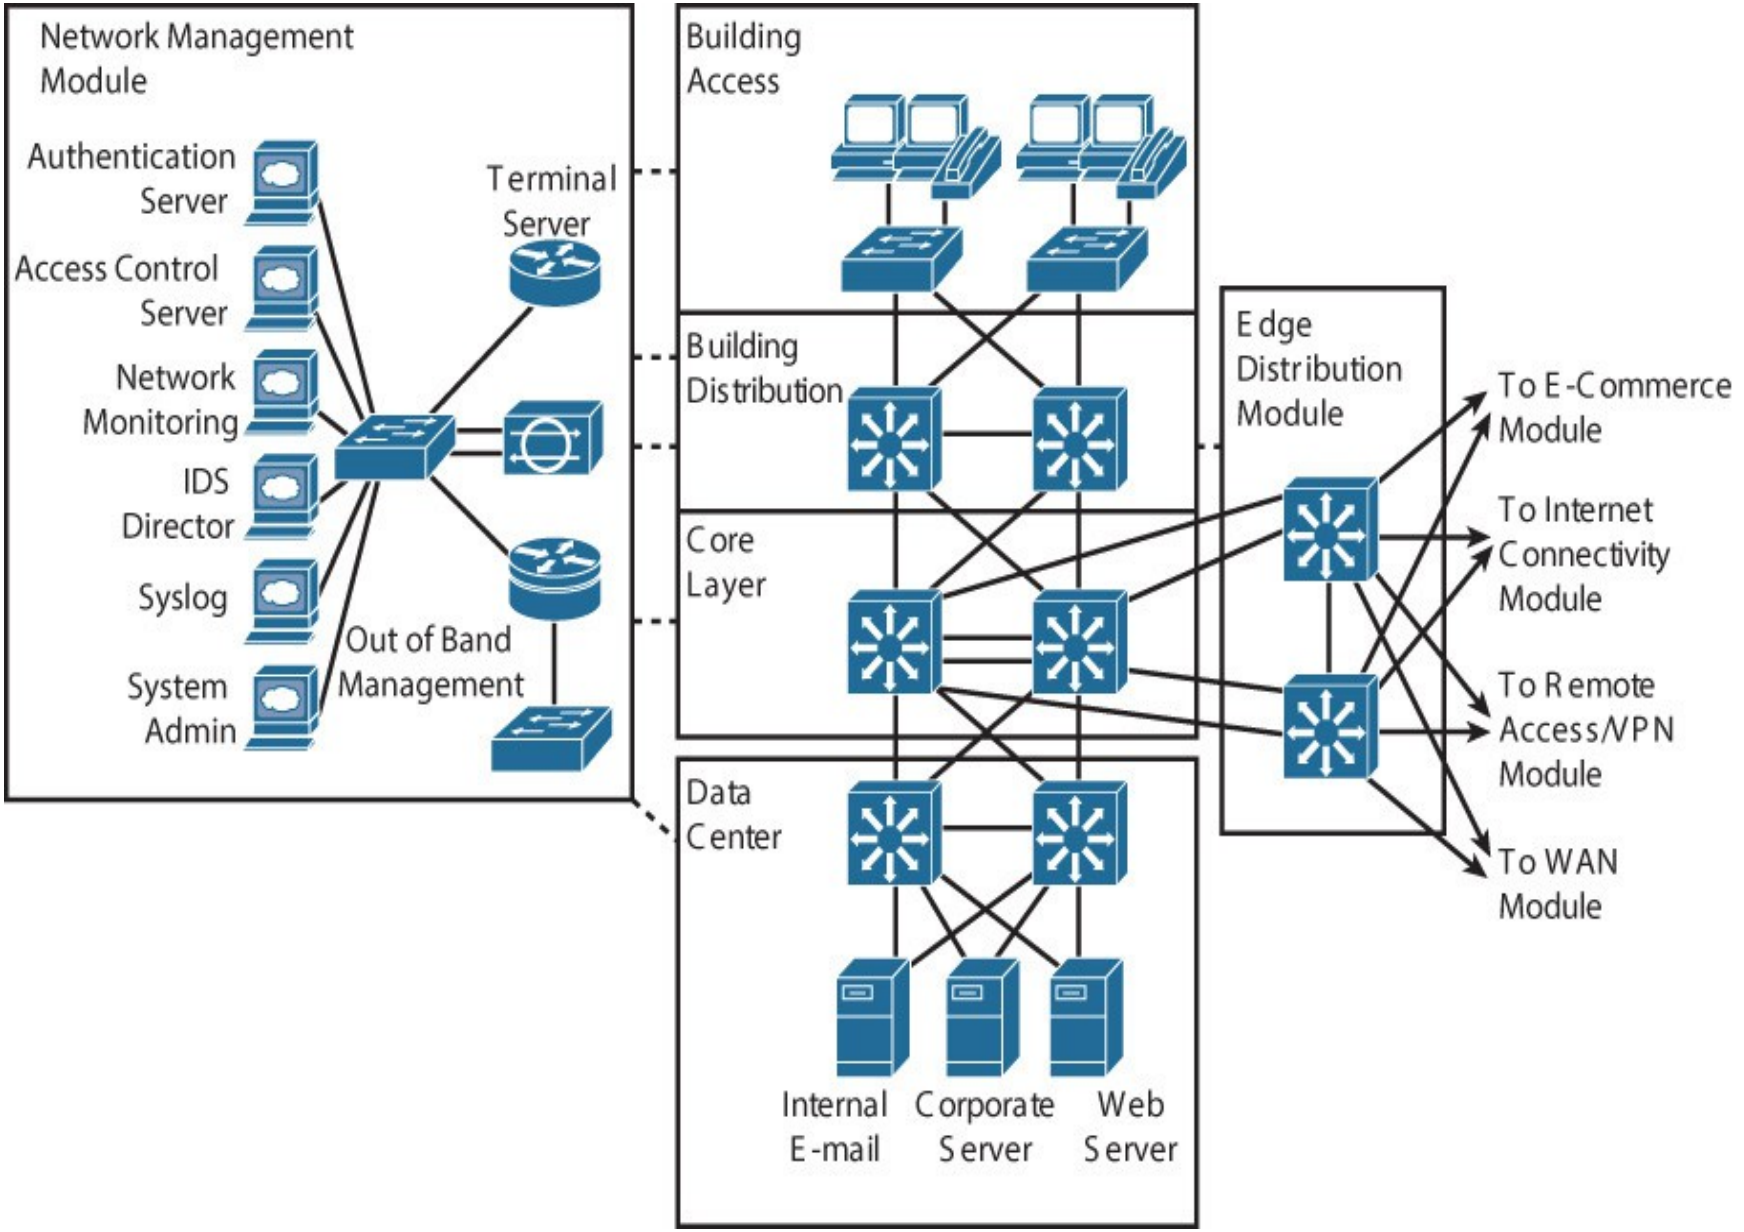

Access Layer

- Provides access for workstations, IP phones, **access points**, printers, video cameras, **and IOT**
- Some of these devices (IP phones, **access points)** extend the network another level
- Formerly pure L2; now L3 is reaching into access
- Connects to distribution layer via redundant links

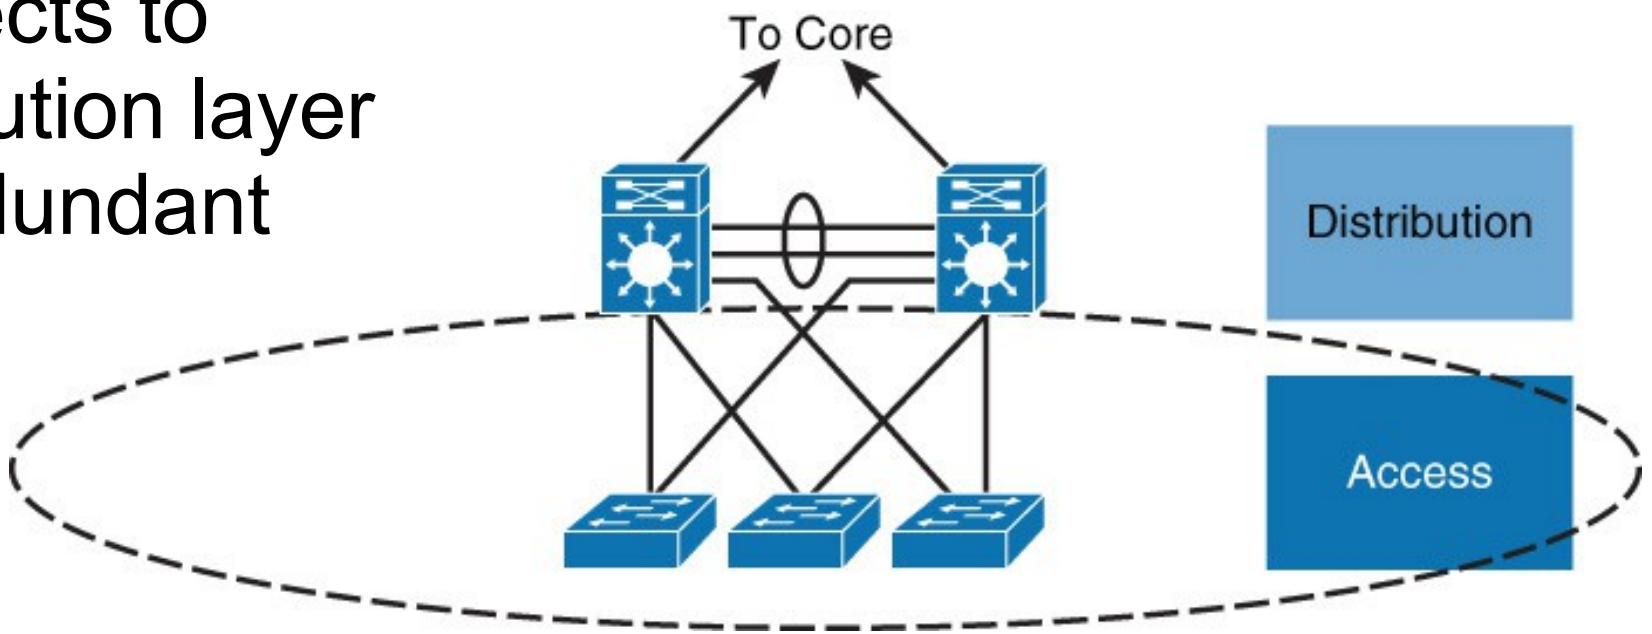

Access Layer Characteristics

"*Provides functionality for end-device connectivity*"

- Previously L2-only due to higher per-port cost for L3 devices; now L3 port costs approach L2 costs
- Redundancy (= HA or High Availability) via STP & multiple uplinks; also HA via FHRP with L3-access
- Convergence: POE for IP phones, Wireless access points, workstation data
- Security: port security, DHCP snooping, DHCP relay, ARP spoofing, IP spoofing, storm control, QoS
- May include tunnels/VPN for remote access

L2 vs L3 in the Access Layer

- Traditional / legacy and still (a little) cheaper than L3
- Allows end-to-end VLANs
- Doesn't scale well (domain size limits)
- Potentially sub-optimal usage of links between access & distribution
- Manageability vs L3 depends on tools used

Only Layer 2 Switching in Access Layer

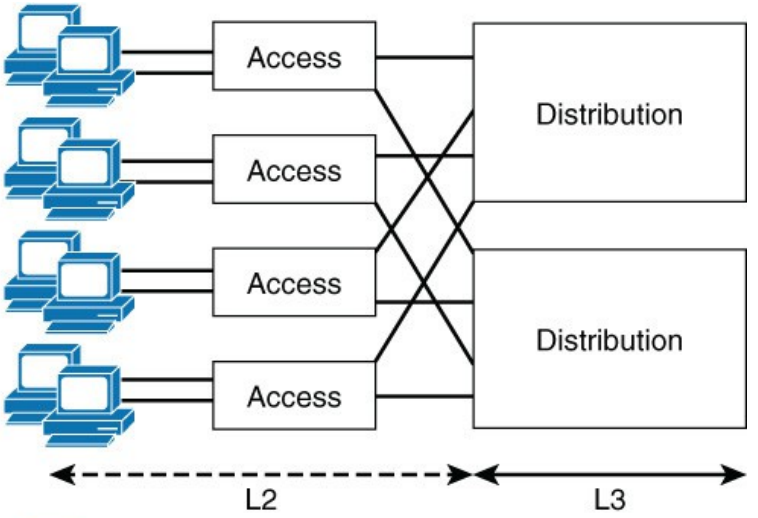

Layer 3 Switching in Access Layer

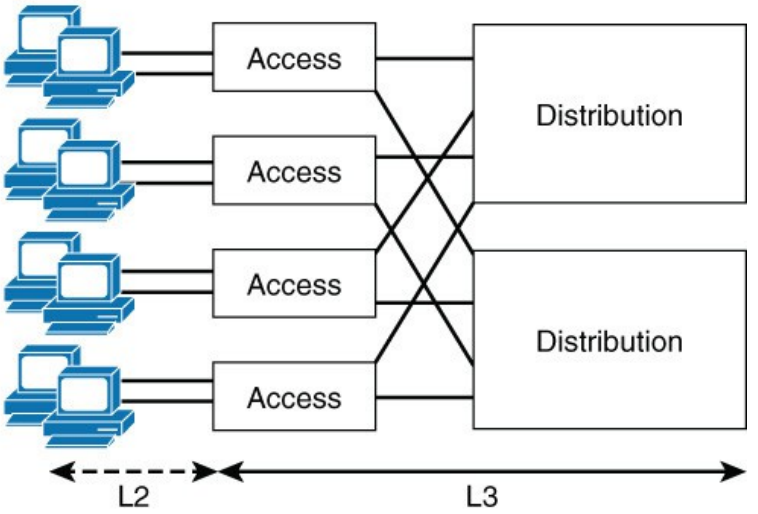

Distribution Layer

- The aggregation layer: access layer switches, wiring closets, WAN connections at the edge of the campus
- May implement policy-based connection (routing); acts as a service and control boundary between access & core (specifics vary with L2/L3 in access)
- Multiple links both up and down

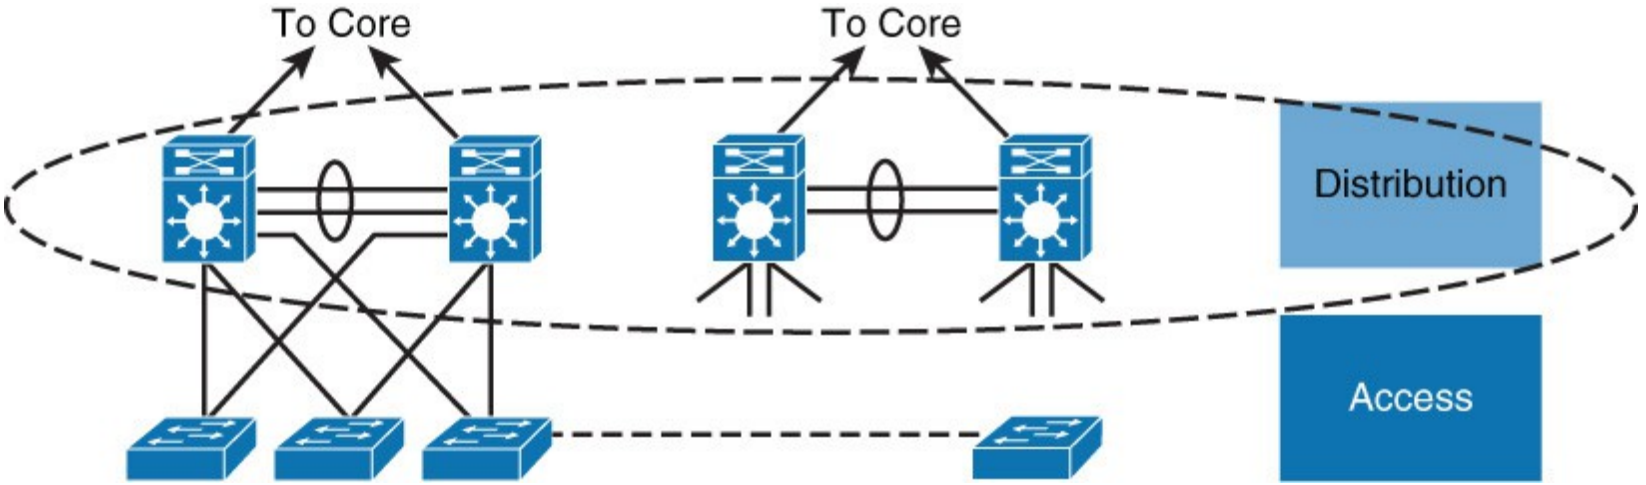

Distribution Layer Characteristics

"*Multi-purpose aggregation & connection layer*"

- In legacy networks, represents the boundary between L2 and L3, terminates VLANs, may be redistribution point between routing domains, and provides HA via FHRP
- With L3-access, represents the boundary between static & dynamic routing: provides a static default route down + summary routes up for access layer, and dynamic routing with the core
- High Availability (HA) via redundant links, fast path recovery, (automatic) load balancing, and bandwidth management via QoS

Core Layer (aka Backbone)

"HA layer that binds all other layers and modules"
Devices are big, dumb, but *fast* bit-movers.

- Requires highest up-time of all layers; Non-Stop Forwarding (NSF) and ability to upgrade-in-service

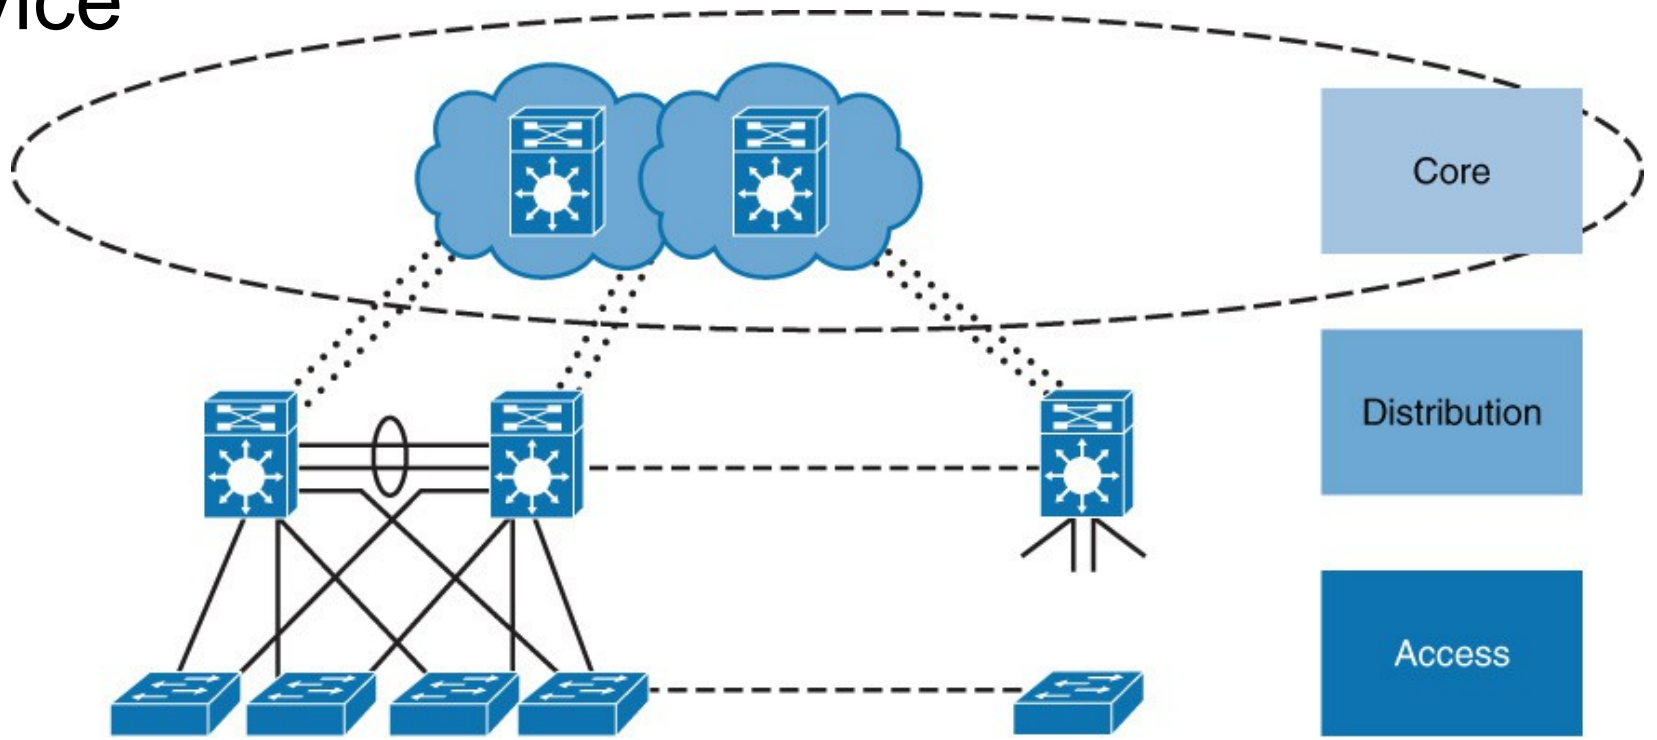

Core Layer

- "Simple" (aka dumb): should not have any complex policy services, routing policies, ACLs, QoS, or connectivity to end-devices
- Has always been a pure L3 layer
- "~~Requires 100% uptime~~" 100% uptime **DNE**!
  - Calculate it: 24 x 3600 = 86400 secs/day
    Even 99.9% uptime provides ~1.5 mins to reboot a (M$) server once a day, every day!
  - Typical best-case figures are "five 9's" or "six 9's" of uptime
- Already covered in Adv. Network Routing, so is not covered in this course

Types of Switches

All network vendors have switches for each of the three layers:

- Access switches are typically "fixed"; 1U or 2U, standard 120/240V AC power
- Distribution switches may be fixed or modular* (fiber-optic vs copper; redundant fans / PSU)
- Core "switches" are usually larger (4-12U),
  - modular, and typically require special power (48+VDC or 120/240VAC @ >15A)
  - may use "Distributed hardware forwarding", meaning forwarding hardware on each line card

*Don't confuse using SFPs as meaning a system is "modular"

L2 Frame Forwarding (store-and-forward)

1. Receive entire frame [input buffer]
2. Re-compute CRC
3. Discard frame if bad CRC; otherwise continue processing
4. Determine VLAN:
   – for access ports, according to configured VLAN
   – for trunk ports, according to VLAN tag on frame (Native VLAN if no tag)
5. Look in MAC address table, according to VLAN, for possible match
   – if match, then select single egress port/trunk
   – if no match, then select flooding as egress method
6. Determine whether to add or strip VLAN tag:
   – Add a tag if passing from access to trunk (new tag = access VLAN #)
   – Strip the tag if passing from trunk to access
   – No change if not crossing access / trunk boundary
7. If VLAN tag added or stripped, re-do CRC
8. Queue the frame on selected egress port(s) [output buffer(s)]

This is "Version 1.0"
(so still not truly complete!)

L2 or L3 Interface Mode

Layer 2: Access or Trunk Ports

Logical Interface (SVI – L3)

Physical Interface (L2 or L3)

```
DLS1# show interface gig 0/2 switchport
Name: Gig0/2
Switchport: Enabled
<output omitted>
```

Enabled ==> L2 mode

- Do interfaces default to Layer 2 or Layer 3?
  - Default on most Access Layer switches: Layer 2
  - Default on Distribution Layer switches: need to check if L2 or L3
- Select L2/L3 mode using the **`switchport`** command:
  `switchport` or **`no switchport`**
- Verify mode:
  `Switch# show interface` {type mod/num} **`switchport`**
  Switchport: Think Layer 2
  Enabled => Layer 2, Disabled => Layer 3

L3 "Switching" Methods

Three generic methods exist for forwarding L3 packets:
– Process switched, Route caching, Topology-based switching
– or as per Cisco terminology: Process, Fast Switching, CEF

- Process switching, or routing table-driven switching:
  – Full lookup is performed at every packet
  – *Control Plane* is responsible for forwarding all packets
- Route Caching or cache-driven switching:
  – First packet always process-switched (*Control Plane*)
  – 2nd+ packets HW switched (*Data plane*) using cache
- Topology-driven switching:
  – all packets forwarded in HW using pre-built FIB table
  – *Data Plane* is responsible for forwarding all packets

Route Caching

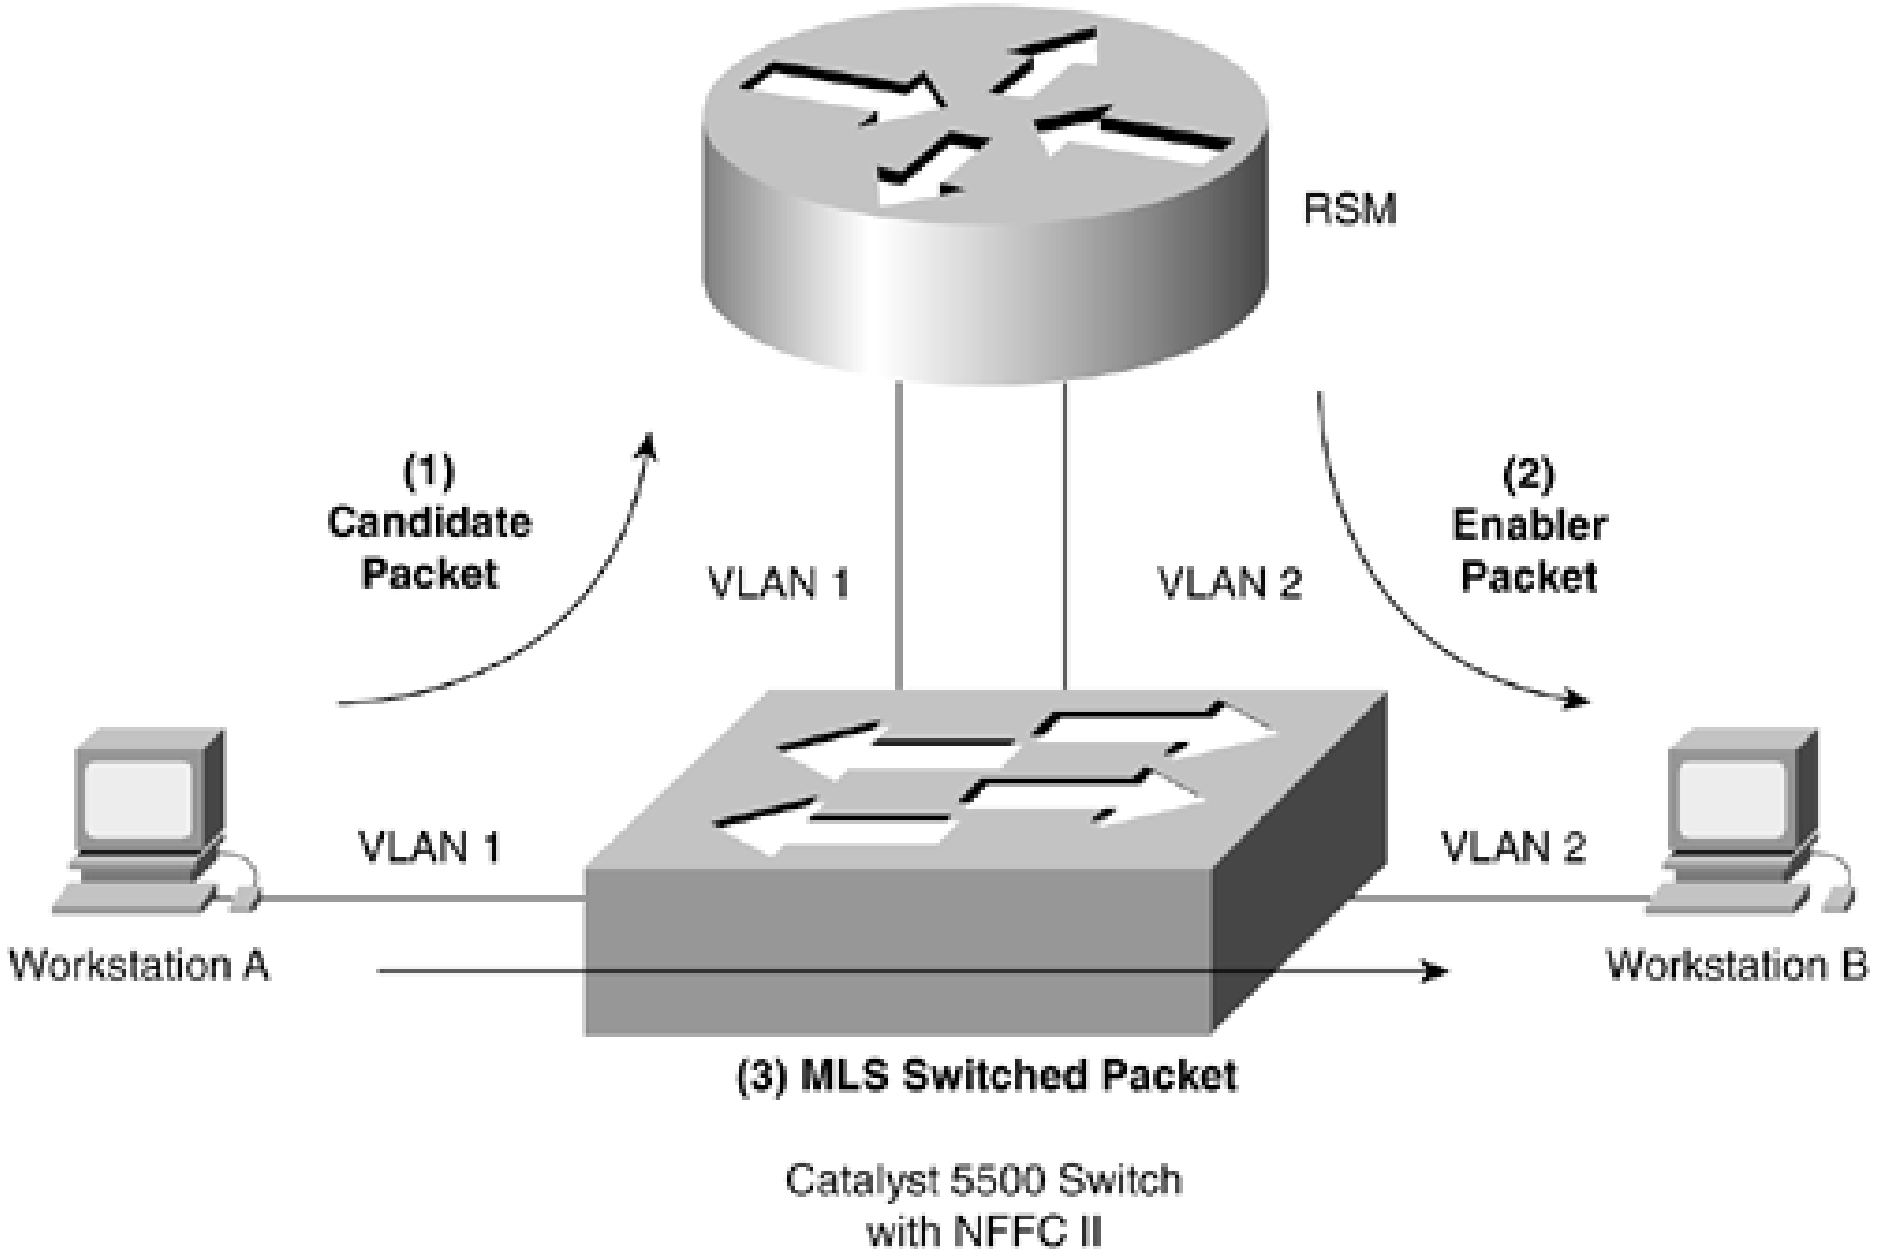

Topology-Based Switching

- Cisco's CEF organizes forwarding via two HW data structures: Forwarding Information Base (FIB) + Layer 2 Adjacency Table
- **FIB**: is essentially a re-ordered copy of the routing table, keeping routes 1:1 along with next-hop information.
  – updated as routing table changes occur
  – high memory/resource usage if routing table is large
- **Adjacency Table**: tracks all nodes reachable via an L2 link and their corresponding encapsulation information
  – built & updated as neighbours are discovered
- Well suited to link-based load balancing
- (Outdated?) CEF may not always enabled by default, even on platforms that support it. To verify, `#show ip cef`
- Can be disabled/enabled globally or per interface.

Topology-Based Switching

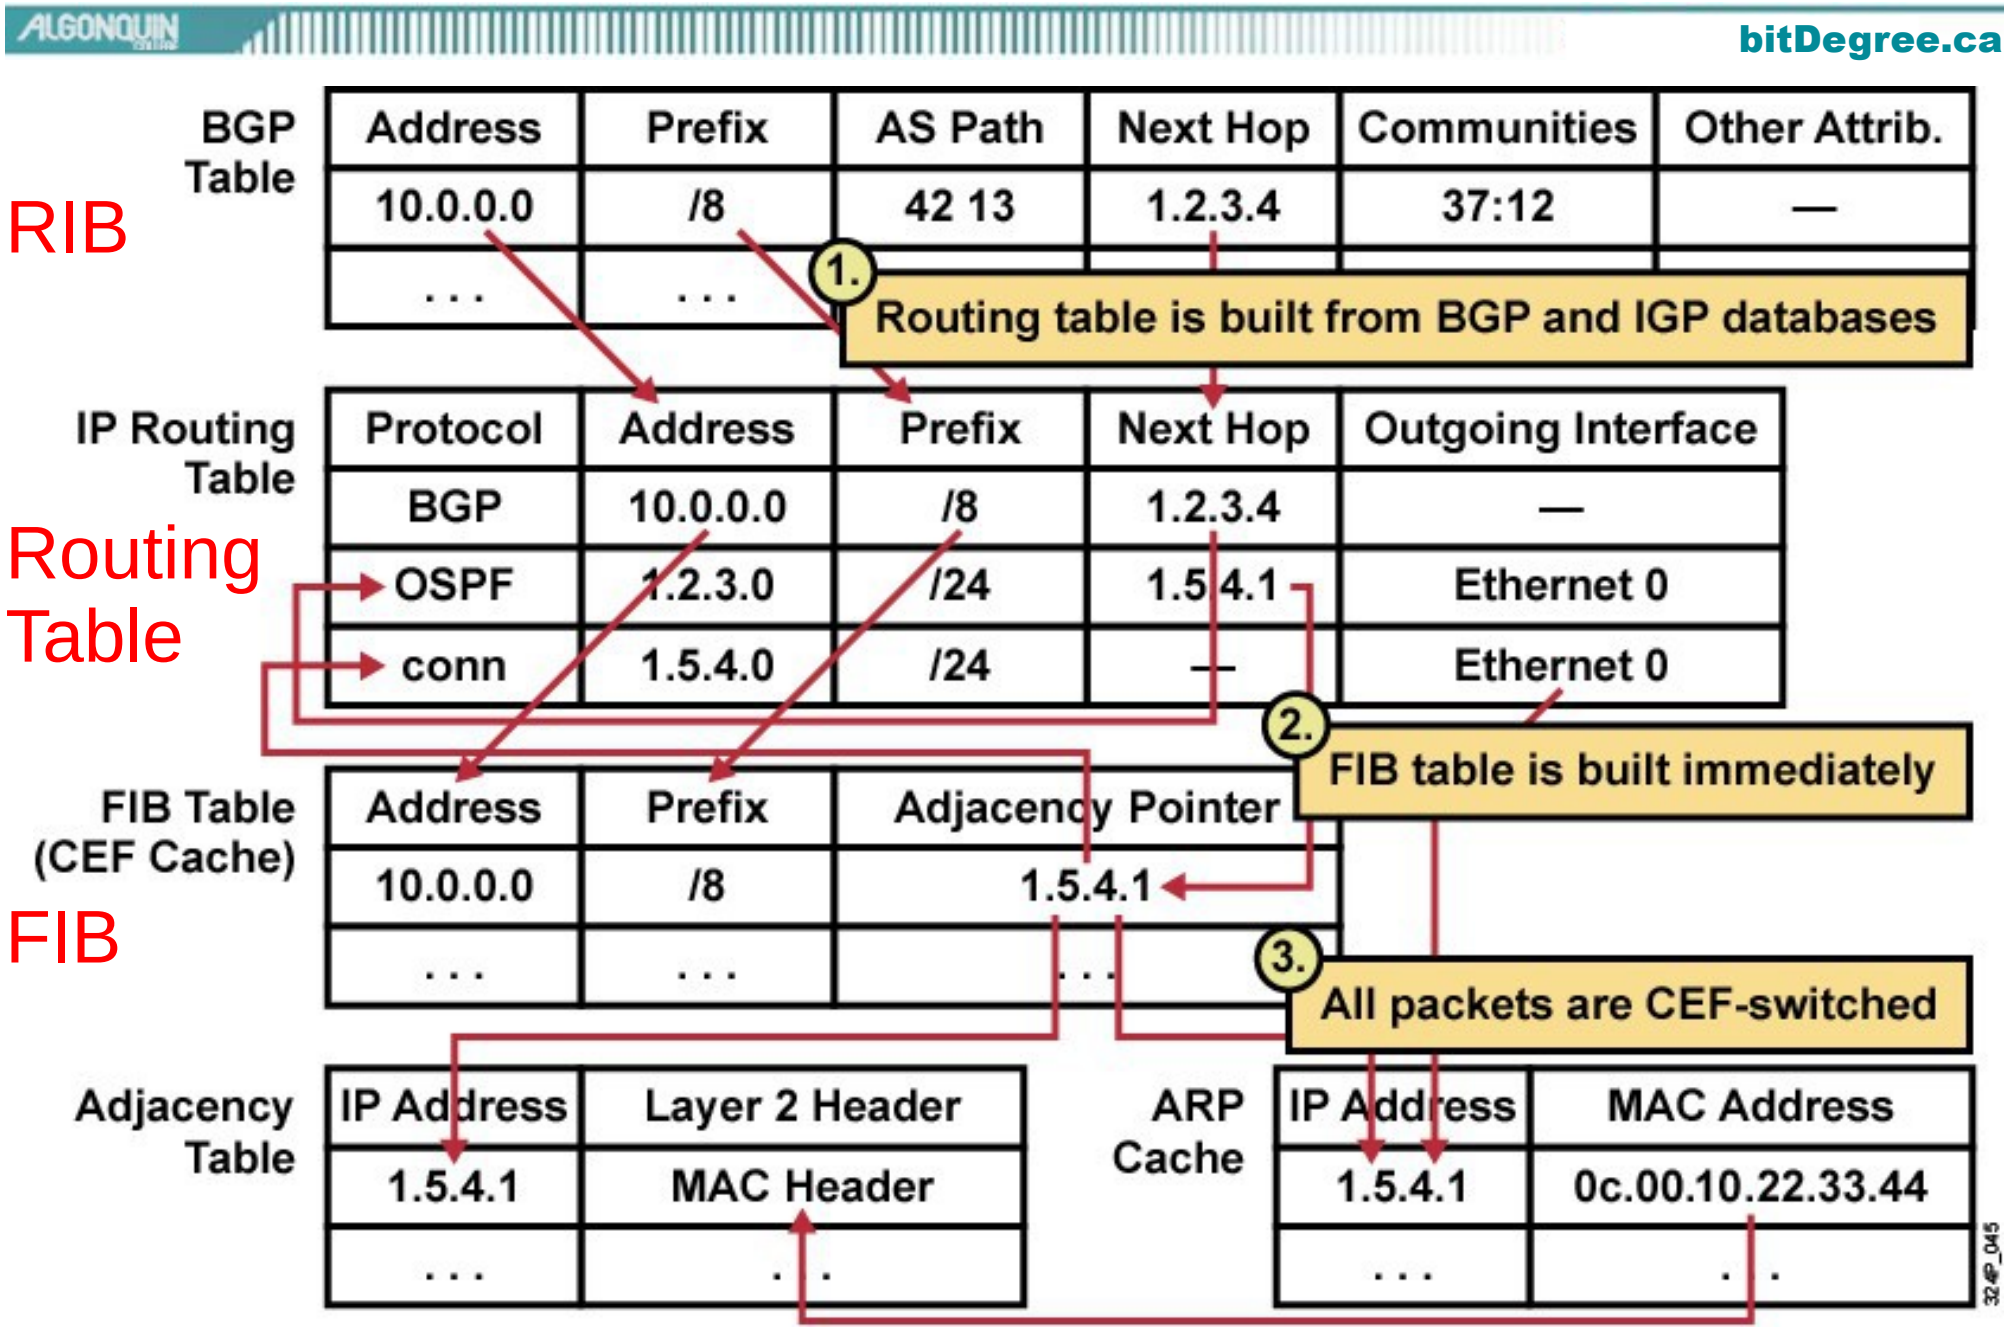

Topology-Based Switching

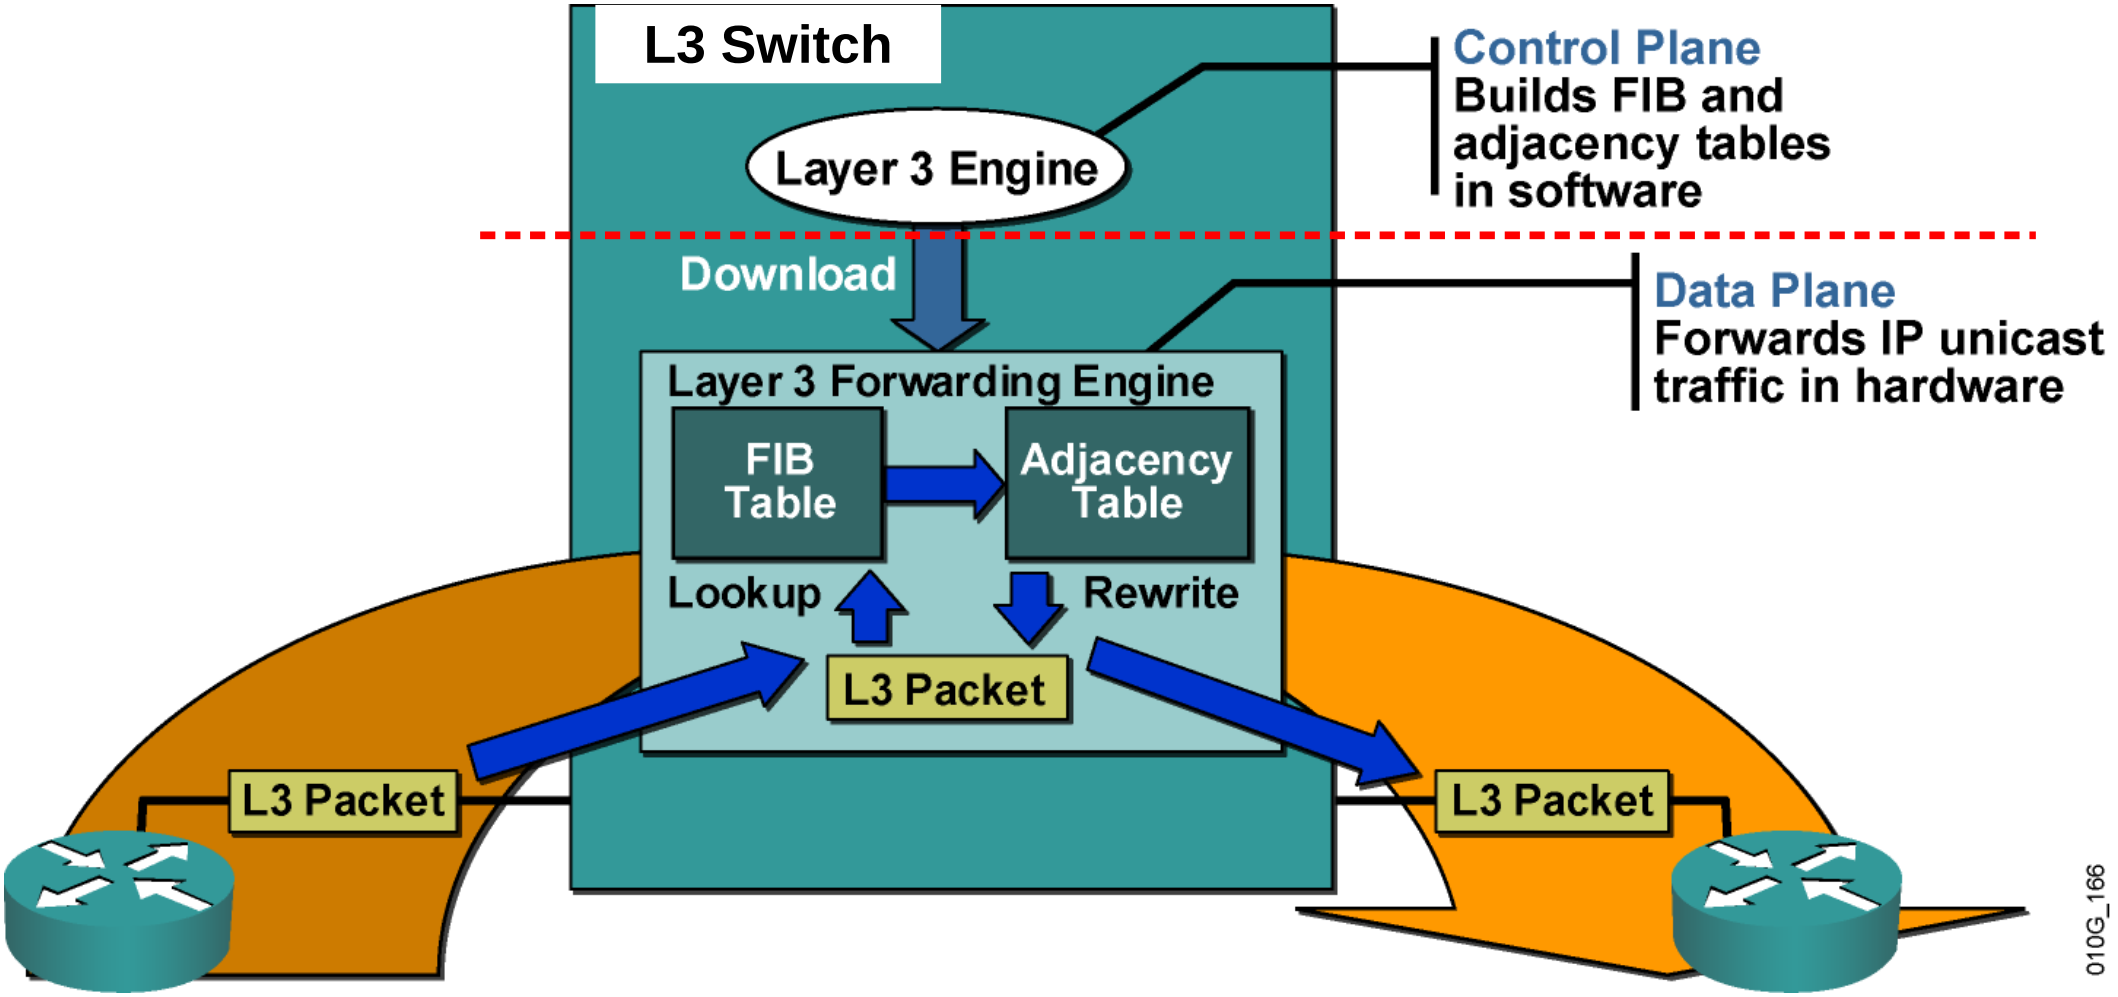

L3 Packet Re-writing

This is "Version 1.0"

1. Check CRC (if bad, drop, done); check dest MAC addr (if not me, drop, done);
   strip L2 framing
2. Check for VLAN tags and strip them (there may possibly be some processing)
   L3 Packet modified / re-written
3. Check IP header details: verify checksum and if okay then check destination IP
   if dest IP is me, punt up to Control Plane, done; else continue forwarding!
4. Decrement TTL count by 1; if now 0, punt up to Control Plane for ICMP msg
   L3 Packet modified / re-written
5. Recalculate Checksum in IP header
   L3 Packet modified / re-written
6. Check routing table for best longest(!) match to destination IP
7. Determine encap required for egress interface: Dest MAC, Src MAC, (VLAN ?)
   L2 frame modified / re-written
8. Compute CRC and add to egress frame
   L2 frame modified / re-written